

~~#19~~
wm

Docket No. 50277-1786 (OID 2001-149-01)

PATENT



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of : Confirmation Number: 4251
Jay Rossiter et al. : Group Art Unit: 2172 **RECEIVED**
Serial No.: 09/945,135 : Examiner: Shahid Al Alam DEC 03 2003
Filed: August 31, 2001 : Technology Center 2100

For: TECHNIQUES FOR MANAGING CONFIGURATION FOR A SYSTEM OF DEVICES ARRANGED IN A NETWORK

APPEAL BRIEF

Mail Stop Appeal Brief – Patents
Hon. Commissioner of Patents
Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is submitted in support of the Notice of Appeal filed June 12, 2003

I. REAL PARTY IN INTEREST

Oracle International Corporation is the real party in interest.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals and interferences.

III. STATUS OF CLAIMS

Claims 1-7 are pending in the application.

Claims 1-7 have been finally rejected in the Final Office Action mailed June 3, 2003. Specifically, Claims 1-7 are rejected under 35 U.S.C. § 102(b) as allegedly being anticipated by U.S. Patent Number 5,606,693 issued to Nilsen et al. ("Nilsen").

It is from this final rejection of Claims 1-7 that this Appeal is taken.

IV. STATUS OF AMENDMENTS

The Claims have not been amended after the Final Office Action.

V. SUMMARY OF THE INVENTION

The magnitude of the expenses associated with maintaining a computer system is due in part to the fact that most computer systems are composed of numerous layers of general purpose components, even though the actual use of all layers is directed to a single specific application. For example, a company may purchase a computer system for the sole purpose of handling the company's accounting. However, the actual workstations purchased for this purpose will typically be general purpose machines that provide all of the hardware services required to be able to handle any application. Many of these services are completely irrelevant for the actual use to which the machine will be directed. For example, the machines may come equipped with built-in "three-dimensional" sound capabilities that will never be utilized while the machines are used for accounting. (Page 2, lines 6-16).

The general purpose machines that are to implement the accounting system will typically come installed with a general purpose operating system. Similar to general purpose hardware, general purpose operating systems attempt to provide services to address all possible needs of all possible types of software. For example, general purpose operating systems employ complex I/O techniques for handling I/O intensive applications, complex techniques to support computation intensive applications, and complex communication techniques to support communication

intensive applications. Many of those services may not be required by the specific context, such as accounting, in which the operating system may actually be used. (Page 2, lines 17-25).

Typically, the hardware and operating system would not be the only "general purpose" components that would be used to implement an accounting system. For example, the accounting system may be configured to interact with a database system. Typically, the database system would provide services to address all possible uses of the database system, and not just those required by the accounting program. Given the diverse environments in which database systems may be used, the services actually required by any given application may be a small fraction of all the services provided by the database system. (Page 2, line 26 – Page 3, line 6).

In the accounting example given above, each of the general purpose components tends to be complex and, consequentially, difficult to manage. When many of such components are thrown together in a single system, the complexity increases exponentially, thus necessitating the significant after-the-purchase computer system maintenance expenses that most companies have encountered. (Page 3, lines 7-11).

In the context of database systems, the complexity of the general purpose database system combined with the complexity of the general purpose operating system with which the database system interacts is typically sufficient to require employment of a full-time database administrator (DBA). It is the responsibility of the DBA to grapple with the complexity of the database system, the operating system, and the interactions between them, so that the other database users can take advantage of the database system without being exposed to underlying complexity. Unfortunately, employing a full-time DBA makes ownership of a useful database system an expensive option. (Page 3, lines 12-19).

Embodiments of the invention provide for a database system that incorporates numerous features that reduce the total cost of maintaining the database system. That database system includes a database appliance that executes a database server on a platform that includes a special purpose operating system specifically tailored to the services required by the database server.

According to certain aspects, the hardware may also be specially tailored to the services required by the database server. According to one aspect of the invention, the database appliance configures itself by detecting the environment in which it resides and setting operational parameters based on the detected environment. The configuration metadata of all components of the system are stored in a centralized repository which itself may reside external to the system. (Page 4, lines 1-11).

According to another aspect, both the database server configuration and the operating system configuration are managed by a remotely located integrated management console, which interacts with and configures the system at the database system level, the operating system level and, according to one embodiment, at the hardware subsystem level. Backup management may also be performed remotely. The remote components, such as the integrated management console, the backup server, and the configuration repository, may communicate with the system through a local area network or wide area network (e.g. the Internet), including through a dial-up connection. (Page 4, lines 12-19).

Using the features described herein, numerous benefits are achieved. Specifically, the ease of using a database system is increased by making the database system easier to set up, manage, and service. The total cost of ownership is reduced by avoiding the costs associated with operating systems and middleware (e.g. web server), by providing remote management, by reducing down time, and by protecting the user's investment. Better performance is achieved through the use of a slim lined operating system, by tuning the database system for a dedicated application, and through the use of custom hardware. The availability of the database system is increased by avoiding problems associated with direct user access to the operating system, and through the use of RAID and/or network attached storage. (Page 4, line 20 – Page 5, line 2).

VI. ISSUES

Whether Claims 1-7 are patentable under 35 U.S.C. § 102(b) as being novel over. *Nilsen*, and

whether the IDS filed on May 29, 2003, Paper No. 12, should be considered if it was filed after the final rejection was mailed.

VII. GROUPING OF CLAIMS

The claims should not be regarded as all standing together since the claims recite respective limitations that render each claim separately patentable. For this appeal, the following groups are recognized:

- A. Independent Claim 1.
- B. Independent Claim 2.
- C. Dependent Claim 3.
- D. Independent Claim 4.
- E. Independent Claim 5.
- F. Independent Claim 6.
- G. Dependent Claim 7.

VIII. ARGUMENT

A. **Independent Claim 1 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.**

Independent Claim 1 is patentable over the prior art because one or more express elements featured in each claim are not disclosed, taught, or suggested by the prior art.

Nilsen is directed towards logging large volumes of data to a plurality of database servers. **Each database server reports status and availability to the configuration controller** which can then adjust future logging requests. The network operator can change the configuration stored in the configuration controller whenever reconfiguration is necessary such as by the addition of new database servers. **A data logging modification is then communicated to each currently active requestor workstation** by the configuration controller (Abstract, emphasis added).

Claim 1 features the elements of:

“gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network;”
“modifying metadata within said centralized repository to initiate configuration changes within said network”
“modifying the operation of one or more of said plurality of devices within said network by propagating said configuration changes from said centralized repository to the devices on said network to cause said configuration changes to take place”

These elements are not disclosed, taught, or suggested by *Nilsen*.

Nilsen lacks any suggestion that the meta data stored in the configuration controller functions as configuration information that **dictates a manner of operation for one or more of a plurality of devices within the network**, as required by Claim 1. In fact, the data stored at the configuration controller does not “dictate the manner of operation” of anything. Rather, the data stored at the configuration controller is merely a collection of data that the database servers have reported to the configuration controller about their status. Thus, while status information flows from the database servers to the configuration controller, the configuration controller does not control anything based on that information. For example, the specification of *Nilsen* makes clear that the configuration data contained within the central configuration controller 132 and 134 merely describes “how many database servers are available and how they are to be accessed” (Col 3, lines 50-52). The configuration controller is thus a mechanism for database servers to communicate to other devices information about how they are working, not a mechanism to communicate to the database servers how they are to work.

Moreover, as the approach of *Nilsen* does not store metadata as claimed within a central repository, *Nilsen* cannot disclose, teach, or suggest the express element of modifying metadata within said centralized repository to initiate configuration changes within said network” as required by Claim 1.

The portion of *Nilsen* cited to show the express element beginning with “modifying the operation” quoted above (Col. 4, lines 2-4) merely states, *in toto*, “As logging proceeds in the database servers, status messages 210 are transmitted over the network 130 from each server to the central configurator 132. An end of data logging message 212 is also transmitted to the

controller 132.” This portion of *Nilsen* merely states that status messages are transmitted from each database server to a central configurator. The status messages do not serve to dictate a manner of operation of the database servers. Further, to the extent that the central configurator communicates information to any other entity, the central configurator merely communicates a message regarding the data logging modification to each currently active requestor workstation. For example, if a database server went off line, the workstations would be notified so data may be logged in a database server that is operational. The significance of this point is that in the approach of *Nilsen*, data flows from the database servers to the central configurator, and then from the central configurator to the requestor workstations; data does not flow from the central configurator to the database servers.

Accordingly, *Nilsen* does not disclose, teach, or suggest the express element of “modifying the operation of one or more of said plurality of devices within said network by propagating said configuration changes from said centralized repository to the devices on said network to cause said configuration changes to take place” because (a) *Nilsen* does not teach modifying the operation of any devices, (b) *Nilsen* does not teach propagating configuration changes from a centralized repository to one or more devices on the network, let alone teaching “propagating said configuration changes from a centralized repository to the device on the network to cause said configuration changes to take place” as required by Claim 1.

Accordingly, as one or more express elements of Claim 1 are absent from *Nilsen*, it is respectfully submitted that Claim 1 is patentable over the cited art and hence are in condition for allowance.

B. Independent Claim 2 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 2 features the elements of:

“gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and
in response to a failure of the system,
recovering the centralized repository from a backup,

using the metadata within the centralized repository to configure the system, and after the system is configured, recovering the system.”

These elements are not disclosed, taught, or suggested by *Nilsen*.

As the first element of Claim 2 is similar to the first element of Claim 1, *Nilsen* does not disclose, teach, or suggest the element of “gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network” of Claim 2 for at least the reasons given above with respect to Claim 1.

Further, *Nilsen* does not disclose, teach, or suggest the element of “in response to a failure of the system, recovering the centralized repository from a backup, using the metadata within the centralized repository to configure the system, and after the system is configured, recovering the system.” First, it is noted that the Office Action neither cites a single portion of *Nilsen* that discloses, teaches, or suggests this element, nor includes a single argument as to why this element is disclosed, taught, or suggested by *Nilsen*. Second, the approach of *Nilsen* lacks the notion of recovering the centralized repository from a backup. The Office Action states that configuration 134 is a backup for configuration 132. While configuration 134 may store the same information as configuration 132, there is no suggestion or discussion in *Nilsen* of recovering configuration 132 from a backup in response to a failure of the system.

Third, the approach of *Nilsen* lacks the notion of using the metadata within the centralized repository to configure the system. Fourth, the approach of *Nilsen* lacks the notion of recovering the system after the system is configured.

Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest one or more elements of Claim 2. Therefore, Claim 2 is patentable over the cited art and hence is in condition for allowance.

C. Dependent Claim 3 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 3 is a dependent claim that directly depends on Claim 1. Claim 3 is therefore allowable for the reasons given above with respect to Claim 1. In addition, Claim 3 introduces one or more additional limitations that independently render it patentable.

In particular, as the approach of *Nilsen* does not store metadata as claimed within a central repository, *Nilsen* cannot disclose, teach, or suggest the express element of “wherein the step of gathering and storing in a centralized repository includes gathering and storing metadata in a centralized repository that resides outside said system” as required by Claim 3. The Examiner relies on elements 132 and 134 of *Nilsen* in rejecting this element; however, elements 132 and 134 do not store metadata as claimed because the specification of *Nilsen* makes clear that the configuration data contained within the central configuration controller 132 and 134 merely describes “how many database servers are available and how they are to be accessed” (Col 3, lines 50-52), as opposed to gathering information that reflects configuration information that dictates a manner of operation for one or more devices.

Further, there is no suggestion in *Nilsen* that metadata is gathered and stored inside a centralized repository that resides “outside said system” as required by Claim 3. Elements 132 and 134 of *Nilsen*, which the Examiner relies upon in showing this element, reside inside the system, rather than “outside the system,” as required by Claim 3.

Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest one or more elements of Claim 3. Therefore, Claim 3 is patentable over the cited art and hence are in condition for allowance.

D. Independent Claim 4 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 4 features the elements of:

“gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network;

managing configuration of said system based upon the metadata within said centralized repository; and
in response to a failure of the system,
configuring the system based on the metadata restored in the centralized repository, and
after the system is configured, recovering the system.”

These elements are not disclosed, taught, or suggested by *Nilsen*.

As the first element of Claim 4 is similar to the first element of Claim 1, *Nilsen* does not disclose, teach, or suggest the element of “gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network” of Claim 4 for at least the reasons given above with respect to Claim 1.

Further, *Nilsen* does not disclose, teach, or suggest the element of “managing configuration of said system based upon the metadata within said centralized repository.” The Office Action neither cites a single portion of *Nilsen* that discloses, teaches, or suggests this element, nor includes a single argument as to why this element is disclosed, taught, or suggested by *Nilsen*, but merely concludes that this express element, which does not appear in either Claim 1 or Claim 2, can be rejected on the same grounds as Claims 1 and 2. However, as previously discussed above, the approach of *Nilsen* does not manage the configuration of the system based upon the metadata within a centralized repository because the data stored within the configurators of *Nilsen* merely describes the status of the database servers, and does not dictate a manner of operation. Consequently, this element is not disclosed, taught, or suggested by *Nilsen*.

Moreover, *Nilsen* does not disclose, teach, or suggest the element of “in response to a failure of the system, configuring the system based on the metadata restored in the centralized repository, and after the system is configured, recovering the system.” The Office Action neither cites a single portion of *Nilsen* that discloses, teaches, or suggests this element, nor includes a single argument as to why this element is disclosed, taught, or suggested by *Nilsen*. The Office Action states that *Nilsen* (at Col. 3, lines 37-38) utilizes a backup/redundant configuration 134 to restore configuration 132. Applicants respectfully disagree with that characterization, as it is submitted that the portion of *Nilsen* cited states only that one controller could become the

primary controller upon failure of the other, and lacks the concept of recovering the controller. However, even if what is asserted by the Office Action were true, “configuring the system based on the metadata restored in the centralized repository, and after the system is configured, recovering the system” as required by Claim 4 is still not shown, because the approach of *Nilsen* does not configure the system based upon metadata stored in a centralized repository. Therefore, it is respectfully submitted that this element is not disclosed, taught, or suggested by *Nilsen*.

Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest one or more elements of Claim 4. Therefore, Claim 4 is patentable over the cited art and hence are in condition for allowance.

E. Independent Claim 5 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 5 features the elements of:

“gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and replicating said system by performing the steps of, copying said metadata to a second centralized repository associated with a second system, and configuring said second system based on the metadata contained in said second centralized repository.”

These elements are not disclosed, taught, or suggested by *Nilsen*.

As the first element of Claims 5 is similar to the first element of Claim 1, *Nilsen* does not disclose, teach, or suggest the element of “gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network” of Claim 5 for at least the reasons given above with respect to Claim 1.

As to the remaining elements of Claim 5, the Office Action rejects the elements for the same reasons associated with Claims 1 and 4, even though the remaining elements of Claim 5 are not present in Claims 1 and 4. The Office Action states that *Nilsen* discloses a redundant

alternative configuration 134 (Col 3, lines 34-36), however the Office Action does not make clear what significance that has. For example, the Office Action earlier characterized configuration 134 as being part of the system, but now appears to characterize configuration 134 as being part of the replicated second system in rejecting Claim 5. It is respectfully submitted that this is an inconsistent position, because in rejecting Claim 4, configuration 134 was characterized as being a backup/redundant component for configuration 132, where configuration 132 and 134 were part of the same system, while in rejecting Claim 5, configuratior 134 is now characterized as part of a second system that was replicated.

Initially, it is respectfully noted that no portion of *Nilsen* suggests that a separate system may be replicated. While *Nilsen* describes using one controller 132 as the primary controller upon failure of the other (Col. 3, 35-38), using one controller in lieu of another is not replicating a system.

Second, no portion of *Nilsen* suggests “copying said metadata to a second centralized repository associated with a second system, and configuring said second system based on the metadata contained in said second centralized repository.” This is so because *Nilsen* does not use metadata as claimed, does not copy metadata to a second centralized repository associated with a second system, and does not configure the second system based on the metadata contained in the second centralized repository. Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest this element.

Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest one or more elements of Claim 5. Therefore, Claim 5 is patentable over the cited art and hence are in condition for allowance.

F. Independent Claims 6 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 6 features the elements of:

“gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and

managing configuration of at least two of an application layer, an operating systems layer, and a hardware layer of said system based upon the metadata within said centralized repository.”

These elements are not disclosed, taught, or suggested by *Nilsen*.

As the first element of Claims 6 is similar to the first element of Claim 1, *Nilsen* does not disclose, teach, or suggest the element of “gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network” of Claim 6 for at least the reasons given above with respect to Claim 1.

Moreover, *Nilsen* does not disclose, teach, or suggest “managing configuration of at least two of an application layer, an operating systems layer, and a hardware layer of said system based upon the metadata within said centralized repository” as required by Claim 6. The Office Action asserts that this element is shown by *Nilsen* because *Nilsen* implicitly discloses two layers by teaching a historical analysis of data (Col. 3, lines 13-15) and the load on each of the database servers (Col. 3, line 62). Assuming, *arguendo*, that this is true, this in no way suggests managing the configuration of these layers based upon metadata within a centralized repository. To the extent that the controllers in the approach of *Nilsen* manage anything, the controllers “provides database server access information to each requesting workstation” (Abstract). Providing information to workstations about the availability and methods of access of database servers in no way suggests the management of a configuration of either an application layer, an operating systems layer, or a hardware layer because those layers cannot be configured merely by communicating to a third party information about their availability and methods of access. Therefore, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest this element.

Consequently, it is respectfully submitted that *Nilsen* does not disclose, teach, or suggest one or more elements of Claim 6. Therefore, Claim 6 is patentable over the cited art and hence are in condition for allowance.

G. Dependent Claim 7 is patentable under 35 U.S.C. § 102(b) as being novel over *Nilsen*.

Claim 7 is directed to a computer readable medium claim that is a multiple dependent claim which directly depends from each of method Claims 1-6. Consequently, Claim 7 is therefore allowable for the reasons given above with respect to Claims 1-6.

H. The IDS filed on May 29, 2003 should be considered.

The IDS filed on May 29, 2003 (hereinafter “the IDS”) should be considered because it is in full compliance with 37 C.F.R. § 1.97 (d). In particular, the IDS is accompanied by (a) the statement specified in paragraph (e) of 37 C.F.R. § 1.97, as evidenced on page 2 of the IDS; and (b) the fee set forth in § 1.17(p), as evidenced by the statement appearing on the IDS on page 3 that “throughout the pendency of this application, please charge any additional fees, including any required extension of time fees, and credit all overpayments to deposit account 50-1302.”

It is acknowledged that the IDS transmittal sheet inadvertently indicated the IDS was filed under 37 C.F.R. § 1.97 (c) instead of 37 C.F.R. § 1.97 (d) due to a clerical error. However, the IDS fulfills the requirements for consideration under 37 C.F.R. § 1.97 (d). Consequently, the Applicants respectfully request that the IDS filed on May 29, 2003 be considered by the Patent and Trademark Office.

CONCLUSION AND PRAYER FOR RELIEF

The rejections under 35 U.S.C. § 102(b) lack the requisite factual and legal basis. The applied reference, *Nilsen*, fails to disclose or suggest numerous claim limitations. Appellants respectfully submit that the imposed rejections under 35 U.S.C. § 102(b) are not viable and respectfully solicit the Honorable Board to reverse each of the imposed rejections under 35 U.S.C. § 102(b) and consider the Information Disclosure Statement of May 29, 2003.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP



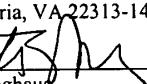
Christopher J. Brokaw
Reg. No. 45,620

1600 Willow Street
San Jose, CA 95125
(408) 414-1080, ext. 201
Date: November 20, 2003
Facsimile: (408) 414-1076

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Brief – Patents, Commissioner for Patents, Box 1450, Alexandria, VA 22313-1450

on November 20, 2003

by 
Sheila Severinghaus

APPENDIX

1. (Previously presented) A method for managing a system that includes a plurality of devices arranged in a network, the method comprising the steps of:
gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network;
modifying metadata within said centralized repository to initiate configuration changes within said network; and
modifying the operation of one or more of said plurality of devices within said network by propagating said configuration changes from said centralized repository to the devices on said network to cause said configuration changes to take place.
2. (Previously presented) A method for managing a system that includes a plurality of devices arranged in a network, the method comprising the steps of:
gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and
in response to a failure of the system,
recovering the centralized repository from a backup,
using the metadata within the centralized repository to configure the system, and
after the system is configured, recovering the system.
3. (Previously presented) The method of Claim 1, wherein the step of gathering and storing in a centralized repository includes gathering and storing metadata in a centralized repository that resides outside said system.
4. (Previously presented) A method for managing a system that includes a plurality of devices arranged in a network, the method comprising the steps of:

gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network;

managing configuration of said system based upon the metadata within said centralized repository; and

in response to a failure of the system,

configuring the system based on the metadata restored in the centralized repository, and

after the system is configured, recovering the system.

5. (Previously presented) A method for managing a system that includes a plurality of devices arranged in a network, the method comprising the steps of:

gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and

replicating said system by performing the steps of,

copying said metadata to a second centralized repository associated with a second system, and

configuring said second system based on the metadata contained in said second centralized repository.

6. (Previously presented) A method for managing a system that includes a plurality of devices arranged in a network, the method comprising the steps of:

gathering and storing in a centralized repository metadata that reflects configuration information about said system, and about each device of said plurality of devices, wherein said configuration information dictates a manner of operation for one or more of said plurality of devices within said network; and

managing configuration of at least two of an application layer, an operating systems layer, and a hardware layer of said system based upon the metadata within said centralized repository.

7. (Previously presented) A computer readable medium carrying one or more sequences of instructions for managing a system that includes a plurality of devices arranged in a network, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of Claims 1, 2, 3, 4, 5, or 6.